

# Commonsense Psychology and the Functional Requirements of Cognitive Models

Andrew S. Gordon

Institute for Creative Technologies  
University of Southern California  
13274 Fiji Way, Marina del Rey, CA 90292 USA  
gordon@ict.usc.edu

## Abstract

In this paper we argue that previous models of cognitive abilities (e.g. memory, analogy) have been constructed to satisfy functional requirements of implicit commonsense psychological theories held by researchers and non-researchers alike. Rather than working to avoid the influence of commonsense psychology in cognitive modeling research, we propose to capitalize on progress in developing formal theories of commonsense psychology to explicitly define the functional requirements of cognitive models. We present a taxonomy of 16 classes of cognitive models that correspond to the representational areas that have been addressed in large-scale inferential theories of commonsense psychology. We consider the functional requirements that can be derived from inferential theories for one of these classes, the processes involved in human memory. We argue that the breadth coverage of commonsense theories can be used to better evaluate the explanatory scope of cognitive models, as well as facilitate the investigation of larger-scale cognitive systems.

## Cognitive Models and Inferential Theories

Computational modeling in the cognitive sciences has come to include two very different research pursuits around the same topic. The first is the pursuit of computational *cognitive models*, software systems that propose testable hypotheses, highlight the inadequacies of current theories, and predict the behavior of people in simulations. For example, an algorithm can be authored to model human similarity judgments and the drawing of analogies (e.g. Falkenhainer et al., 1989) and used to inform a theory about how these cognitive functions are computed in people (e.g. Gentner & Markman, 1997). The second pursuit is the development of *inferential theories*, software systems that propose representations and inference mechanisms that describe the explanations and predictions that people generate. When these explanations and predictions are themselves about human cognition, it falls under the heading *commonsense psychology*, also referred to as naïve psychological reasoning. For example, a logical theory of human emotions (e.g. Sanders, 1989) can be paired with an automated reasoning engine to generate inferences about human emotional states that

match those that people would make. While there is certainly some very interesting relationship between these two pursuits, they have very different aims. Cognitive models are authored to describe the way people think (the processes of human cognition). Inferential theories about the mind are authored to describe the way people think they think (the inferences that people make about human cognition).

The interesting relationship between these two pursuits has not yet been explored within the cognitive sciences. The reason for this is largely due to the amount of debate that has arisen around one area of human cognition where these two pursuits are in contention: cognitive models of Theory of Mind reasoning. Originally investigated as an ability that young children acquire to reason about the false beliefs of other people (Wimmer & Perner, 1983), the study of Theory of Mind reasoning has come to include a range of social cognition behaviors, perspective taking, metacognition, and introspection (Baron-Cohen et al., 2000). Two competing theories of Theory of Mind reasoning have been proposed. Proponents of Theory Theory have argued that Theory of Mind reasoning relies on tacit inferential theories about mental states and processes (inferential theories of commonsense psychology), which are manipulated using more general inferential mechanisms (Gopnick & Meltzoff, 1997; Nichols & Stich, 2002). Proponents of Simulation Theory argue that Theory of Mind reasoning can be better described as a specialized mode of reasoning where inferences are generated by employing one's own reasoning functions (which can be described as cognitive models) to simulate the mental states and processes of other people (Goldman, 2000).

Resolution of this debate in the future may, indeed, show that it is more productive to understand Theory of Mind reasoning as employing either inferential theories of the mind or cognitive processes in simulation mode. However, the victory of either of these positions in this particular debate will not invalidate either pursuit. That is, the success of Simulation Theory would not conclude that people do not have and employ inferential theories of commonsense psychology, just as the success of Theory Theory would not imply that people lack cognitive

<b>Report Documentation Page</b>			<i>Form Approved OMB No. 0704-0188</i>		
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>					
1. REPORT DATE <b>2006</b>	2. REPORT TYPE	3. DATES COVERED <b>00-00-2006 to 00-00-2006</b>			
4. TITLE AND SUBTITLE <b>Commonsense Psychology and the Functional Requirements of Cognitive Models</b>			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>University of California, Institute for Creative Technologies, 13274 Fiji Way, Marina del Rey, CA, 90292</b>			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES <b>7</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

processes other than inference. By focusing on the incompatibilities of these two pursuits in modeling Theory of Mind reasoning abilities, we are ignoring the interesting ways that these pursuits overlap, and how progress in one can inform progress in the other.

In this paper, we focus specifically on how progress in the development of inferential theories of the mind can be used to guide the development of integrated, large-scale cognitive models of heterogeneous cognitive processes. For the most part, we ignore the debate surrounding Theory of Mind reasoning, and investigate instead the relationship between mental processes that are modeled in inferential theories and the functional requirements of computational cognitive models. First, we describe a large-scale effort to author commonsense psychological theories that achieves both broad coverage and inferential competency. Second, we outline the suite of cognitive models that are implied by this work, presented as a list of 16 functionally delineated classes. Third, we look specifically at how logical theories of cognitive processes can be used to specify the functional requirements of computational cognitive models, using as an example the functions of a cognitive model of human memory. We conclude with a discussion of future work.

## Logical Theories of Commonsense Psychology

Interest in developing inferential theories of the mind has been persistent throughout the history of research on formal commonsense knowledge representation, but has typically been secondary to efforts aimed at developing inferential theories of physical systems (see Davis, 1998). Progress in this area has often been made by focusing on narrow aspects of human reasoning (e.g. Cohen & Levesque, 1990) or by adopting non-formal notation (Smedslund, 1997). It has only been recently that large-scale formalizations of commonsense psychological knowledge have been attempted that achieve both breadth of coverage and inferential competency. Here we highlight the work of Gordon (2004) and Hobbs & Gordon (2005). In an effort to identify the representational requirements of human strategic planning, Gordon (2004) conducted a large-scale analysis of 372 planning strategies gathered from sources in 10 different real-world domains. This analytic approach involved authoring pre-formal representations of each planning strategy with the aim of identifying each of the concepts that would have to be formalized in order to correctly define the strategy across analogous planning cases. Of the 988 unique concepts that were identified in this work, two-thirds dealt with the mental states and processes of people. Gordon (2002) organized this subset of concepts into 30 representational areas (e.g. human memory, emotions, plan following), a comprehensive characterization of the breadth of concepts involved in human commonsense psychological reasoning.

Hobbs & Gordon (2005) then began an effort to develop inferential theories based on these 30 representational areas to support automated commonsense inference. The aim of

this work is to develop formal (logical) theories that achieve a high degree of coverage over the concepts related to mental states and processes, but that also have the necessary inferential competency to support automated commonsense reasoning in this domain. These formal theories were authored as sets of axioms in first-order predicate calculus, enabling their use in existing automated reasoning systems (e.g. resolution theorem-proving algorithms).

Gordon & Hobbs (2003) give an example of inferential theories that are being authored using this approach, namely a commonsense theory of human memory. The topic of human memory concerns memories in the minds of people, which are operated upon by memory processes of storage, retrieval, memorization, reminding, and repression, among others. The formal theory of commonsense human memory presented by Gordon & Hobbs supports inference about these processes with encodings of roughly three-dozen memory axioms in first-order predicate calculus. Key aspects of this theory can be characterized as follows:

1. Concepts in memory: People have minds with at least two parts, one where concepts are stored in memory and a second where concepts can be in the focus of one's attention. Storage and retrieval involve moving concepts from one part to the other.

2. Accessibility: Concepts that are in memory have varying degrees of accessibility, and there is some threshold of accessibility for concepts beyond which they cannot be retrieved into the focus of attention.

3. Associations: Concepts that are in memory may be associated with one another, and having a concept in the focus of attention increases the accessibility of the concepts with which it is associated.

4. Trying and succeeding: People can attempt mental actions (e.g. retrieving), but these actions may fail or be successful.

5. Remember and forget: Remembering can be defined as succeeding in retrieving a concept from memory, while forgetting is when a concept becomes inaccessible.

6. Remembering to do: A precondition for executing actions in a plan at a particular time is that a person remembers to do it, retrieving the action from memory before its execution.

7. Repressing: People repress concepts that they find unpleasant, causing these concepts to become inaccessible.

Swanson & Gordon (2005) describe an effort to evaluate the inferential competency of Gordon & Hobbs' theory of human memory. The approach attempts to prove (using automated theorem-proving techniques) the validity of a commonsense strategy that people use to manipulate their own memory processes by focusing on concepts that are associated to stored memories. This work highlights some of the inferential limitations of relying solely on logical theories to support Theory of Mind reasoning. Using the theory, it is possible to infer that people are more likely to remember something when they have an associated concept in their focus of attention. For example, the

memory of turning on the dishwasher might be inspired when thinking about some dishwashing soap. However, a prediction as to whether they would actually remember is dependent on how accessible the memory was in the first place, and how strong the association between the two concepts actually is. This is information that would be unlikely encoded as part of a inferential theory of the mind (Theory Theory), but could be obtained by deferring to one's own memory processes in simulation mode (Simulation Theory).

## A Taxonomy of Cognitive Models

The 30 representational areas that Hobbs & Gordon (2005) formalize as inferential theories have a high degree of overlap (in name, at least) with the classes of cognitive models that have historically been developed. To illustrate this idea more completely, consider the effect of treating these 30 areas not as a taxonomy of inferential theories that need to be authored, but rather as a taxonomy of cognitive models that would participate in an integrated cognitive architecture.

Below is a list of 16 functional classes of cognitive models, where the original 30 representational areas are regrouped according to function (rather than representational) distinctions.

1. Knowledge and inference (Managing knowledge): Models of how people maintain and update their beliefs in the face of new information (e.g. Byrne & Walsh, 2002).
2. Similarity judgment: Models of how people judge things to be similar, different, or analogous (e.g. Gentner & Markman, 1997).
3. Memory: Models of memory storage and retrieval (see Conway, 1997).
4. Emotion: Models of emotional appraisal and coping strategies (e.g. Gratch & Marsella, 2004).
5. Envisionment (including Execution envisionment): Models of how people reason about causality, possibility, and intervention in real and imagined worlds (e.g. Sloman & Lagnado, 2005).
6. Explanation (including Causes of failure): Models of the process of generating explanations for events and states with unknown causes (e.g. Leake, 1995).
7. Expectation: Models of people come to expect that certain events and states will occur in the future, and how they handle expectation violations (e.g. Schank, 1982).
8. Theory of Mind reasoning (Other agent reasoning): Models of how people reason about the mental states and processes of other people and themselves.
9. Threat detection: Models of how people identify threats and opportunities that may impact the achievement of their goals (e.g. Pryor & Collins, 1992).
10. Goal management (including Goals and Goal themes). Models of how people prioritize and reconsider the goals that they choose to pursue (e.g. Schut et al., 2004).
11. Planning (Plans, Plan elements, Planning modalities, Planning goals, Plan construction, and Plan adaptation):

Models of the process of selecting a course of action that will achieve one's goals (e.g. Rattermann, 2001).

12. Design: Models of how people develop plans for the creation or configuration of an artifact, process or information.

13. Scheduling: Models of how people reason about time and select when they will do the plans that they intend to do.

14. Decision making: Models of how people identify choices and make decisions (e.g. Zachary et al., 1998).

15. Monitoring: Models of how people divide their attention in ways that enable them to wait for, check for, and react to events in the world and in their minds (e.g. Atkin & Cohen, 1996).

16. Plan execution (Execution modalities, Repetitive execution, Body Interaction, Plan following, Observation of execution): Models of the way that people put their plans into action and control their own behavior (e.g. Stein, 1997).

While this taxonomy does not fully catalogue the entire breadth of cognitive modeling work that has been done (e.g. for language acquisition), it is interesting that a large amount of previous modeling work can be neatly categorized into commonsense psychological divisions. This observation is further supported in that few previous cognitive modeling efforts have been attempted for cognitive functions that cross several commonsense psychological areas. This is not at all coincidental, as these cognitive models have been constructed to satisfy functional requirements of implicit commonsense psychological theories held by researchers and non-researchers alike. For example, when researchers work on cognitive models of memory, they rely on their own commonsense view of human memory to help them determine what cognitive behaviors their models should explain. The potential pitfall of this approach is that human cognition may be best described using vocabulary and functional delineations that are far removed from our commonsense psychological views (e.g. a neural process account). However, it is difficult to imagine how descriptions that are not related to our commonsense views of psychological processes can be accepted as satisfactory explanations.

Rather than working to avoid the influence of commonsense psychology in cognitive modeling research, the opposite approach may be more productive. By investigating human commonsense psychological theories more thoroughly and explicitly, we can better describe the functional properties of cognitive process that cognitive models should explain. That is, if our commonsense theories of a cognitive process include certain states, events, and controls, then the computational cognitive models of these processes should define (where possible) these states, events, and controls in the algorithms and data structures of the model.

Practically speaking, this approach requires that highly descriptive characterizations of states, events, and controls of commonsense psychological processes are available to

serve as a design specification. This is exactly the place where current progress on authoring formal inferential theories of commonsense psychology can be applied to cognitive modeling. Although these models are being authored using formal notation rather than functional notation, the translation from one to the other is straightforward. In general, the resulting functional specifications will include a broader range of states, events, and controls than have been attempted in previous cognitive modeling efforts, and can be used to guide progress in the future. To support this claim, the next section elaborates the functional requirements of cognitive models of human memory, as identified by a formal treatment encoded in an inferential theory.

## Functional Requirements of Memory Models

The study of human memory has an enormously rich history in cognitive psychology (e.g. Bartlett, 1932). Many different approaches have been pursued to develop an understanding of memory processes, including the authoring of computational cognitive models. A good example of the utility of computational cognitive models is seen in the Forbus et al. (1994) model of similarity-based memory retrieval. In their two-stage model, a target situation in working memory serves as a retrieval cue for a possible base situation in long-term memory. In the first stage, a fast comparison process is done between a target and potential bases using a flat feature-vector representation, resulting in a number of candidate retrievals. In the second stage, attempts are made to identify deep structural alignments between the target and these candidates using a graph-comparison algorithm. Based on the strength of the comparisons made in these two stages, base situations that exceed a threshold are retrieved. This computational model has helped explain the empirical evidence of human memory retrieval performance, including why reminders are sometimes based only on surface-level similarities, and other times based only on deep structural analogies.

Part of the elegance of this model of memory retrieval is in its functional simplicity. The system is initialized with a database of situations to be stored in long-term memory. Its processes are initiated when a target situation is in working memory. Its sole effect on other cognitive processes is the retrieval of base situations from long-term memory into working memory.

Gordon & Hobbs' (2003) formal inferential theory, described earlier, is not a cognitive model. Instead of explaining how human memory works, it encodes a commonsense view of how people think human memory works. As such, it describes a human memory mechanism with significantly more functionality than outlined in Forbus et al. (1993). To illustrate this point, consider the following functional translations that can be made from the logical forms presented in the inferential theory.

First, there are the two methods for storing and retrieving information from memory. The storage method can be encoded as a procedure (having no return value)

that is called with one argument, namely the information that is to be stored into memory. The retrieval method can be encoded as a function (returning the information from memory), which takes as an argument the information that is used as a retrieval cue. The retrieval function can fail in the case that the information is not accessible (i.e. it throws a retrieval failure exception).

```

procedure: Memory storage
  arguments: <Information to be stored>
function: Memory retrieval
  arguments: <Cue information>
  returns: <Information from memory>
  error: <Memory retrieval failure>

```

Importantly, the memory retrieval function above does not correspond to the functionality supported in the model of Forbus et al., where situations are retrieved in reaction to the presence of a different situation in working memory. Instead, the memory retrieval function above references the act of trying to remember some specific information (e.g. someone's name) given a retrieval cue (e.g. their spouse's name). To handle the memory-initiated case described in the cognitive model, Gordon & Hobbs formalize the case of having information in one's focus of attention (i.e. a situation in working memory), and being reminded of associated information.

```

procedure: Focus on
  arguments: <Information in the focus of
             attention>
event: Reminding
  returns: <Information from memory into the
          focus of attention>

```

The next set of functions translated from the inferential theory push further away from the functionality supported by the cognitive model. In the commonsense view, one of the key purposes of memory is to make sure that people remember to do things that they intend to do, at the right time. That is, people schedule to execute actions at times in the future (encoded as a procedure), and it is memory's job to retrieve these actions into the focus of attention at that future time. As in the case of reminding, remembering to do can be encoded as the event of its occurrence (to be processed by an event-driven control loop). Here forgetting to do can also be encoded as an event, when a person retrieves the actions that are to be taken, but where the time to do them has already past.

```

procedure: Schedule to do
  arguments: <Action to take at some time>
event: Remember to do
  returns: <Action to take at the current time>
event: Forget to do
  returns: <Action to take at some previous
          time>

```

The final capability that is expressed the commonsense model of memory, and not supported by the cognitive model, concerns the phenomenon of memory repression.

Although the cognitive validity of memory repression continues to be debated (see Anderson & Green, 2001), this capacity has been a part of our commonsense models of memory at least since the popularization of the work of Sigmund Freud. Functionally, memory repression can be viewed as a procedure that is called of the memory mechanism and that takes as an argument the information to be repressed.

*procedure: Repress information  
arguments: <Information to repress from  
memory>*

The cognitive model of Forbus et al., like many other computational cognitive models of memory that have been proposed, does not provide an explanation for each of these commonsense memory functions. It is certainly the case the cognitive models could be elaborated to support these functions. For example, memory storage could be handled in the Forbus et al. model as a simple addition of a situation into long term memory, while memory retrieval could be handled by authoring a separate second-stage algorithm further evaluated the appropriateness of the base given the retrieval cue (a type of target).

The point here is not to criticize the quality of the computational cognitive models that have been proposed in the past, but rather to explicitly identify the corresponding commonsense psychological functions that these models are attempting to explain. The primary means of evaluating the quality of these models remains the same: how predictive they are of the empirical evidence. However, we can now begin to evaluate cognitive modeling efforts by secondary criteria, namely how much of the breadth of commonsense mental phenomenon the model explain. Progress in the development of large-coverage inferential theories of commonsense psychology supports this secondary evaluation by explicitly identifying the commonsense functional requirements of cognitive models.

## Discussion

Forbus (2001), in reflecting on progress made in the development of computational cognitive models of analogy, argued that further advancement in this area requires that modeling efforts scale up to explore the role of analogy “in the large”, i.e. in relation to other larger-scale cognitive processes. Forbus offered a new constraint that should be placed on cognitive models, namely the integration constraint: they should be usable as a component of a larger cognitive system. Forbus et al. (2002) further pursued this aim by describing the functions and structures of a cognitive model of analogy as predicate calculus formulas, enabling the integration of analogical reasoning with formal inference mechanisms used in Artificial Intelligence.

The work that is described in this paper also aims to support the integration of cognitive models into larger reasoning systems, but we take exactly the opposite

approach. Rather than deriving the representations that are used in inferential systems from the functionality supported by current cognitive models, we define the functionality that should be supported by cognitive models by translating the formal semantics of inferential theories. While both approaches recognize that making the connections between inferential theories and cognitive models can facilitate integration, the question concerns which of these two should be viewed as the most appropriate starting point. In this paper we have argued that current commonsense psychological theories (e.g. Hobbs & Gordon, 2005) have the advantage of breadth of coverage.

Defining the functional requirements of cognitive models based on commonsense psychological theories facilitates larger-scale cognitive modeling precisely because an intuitive understanding of these functions is already shared among members of disparate research groups. As an engineering enterprise, the development of integrated, larger-scale cognitive models by distributed sets of researchers is facilitated if both the delineation of modules (e.g. memory, emotion, planning) and functions within these modules (e.g. retrieval, reminding, repression) are in accordance with one’s own commonsense view. This will be particularly important in supporting some degree of computational modularity, where competing computational cognitive models in larger systems can be directly swapped for one another in order to conduct comparative evaluations.

In addition to the engineering benefit of this approach, explicating the connection between inferential theories of commonsense psychological and human cognitive function opens the door for investigations of human reasoning where these two components interact. Perhaps the best example of this concerns the production and interpretation of language that references mental processes. Research in this area (e.g. Gordon et al., 2003) has made progress by assuming that the semantics of linguistic expressions related to mental states and processes can be defined entirely within the scope of inferential theories of commonsense psychology. However, a deeper understanding of the relationship between commonsense psychological theories and cognitive processes will be necessary to account for cases where the language has a direct effect on cognitive function (e.g. imperatives of mental actions).

Ultimately, progress in understanding the relationship between commonsense psychological theories and the functions implemented in computational cognitive models will enable us to better tackle the debate concerning Theory of Mind reasoning. It is only through the parallel development of inferential theories and cognitive models that we can appropriately assess the limitations of each and determine where the real differences in these two approaches actually exist. It is in this debate that the real benefits of computational modeling (both types) are evident, providing a degree of formality and clarity that is otherwise difficult to achieve. For now, a resolution of this

debate will have to wait as further progress is made, and both approaches find ways to capitalize on the strengths of the other.

### Acknowledgments

The project or effort depicted was or is sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM), and that the content or information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

### References

Anderson, M. & Green, C. 2001. Suppressing unwanted memories by executive control. *Nature* 410, 366-369.

Atkin, M. & Cohen, P. 1996. Monitoring strategies for embedded agents: Experiments and analysis. *Journal of Adaptive Behavior* 4(2), 125-172.

Baron-Cohen, S., Tager-Flusberg, H. & Cohen, D. (Eds.) 2000. *Understanding Other Minds: Perspectives from Developmental Cognitive Neuroscience*, 2nd Edition. New York: Oxford University Press.

Bartlett, F. 1932. *Remembering: A study in experimental and social psychology*. Cambridge, England: Cambridge.

Byrne, R. & Walsh, C. 2002. Contradictions and counterfactuals: Generating belief revisions in conditional inference. *Proceedings of the 24th Annual Meeting of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.

Cohen, P. & Levesque, H. 1990. Intention is choice with commitment. *Artificial Intelligence* 42(3), 213-261.

Conway, M. (ed) 1997. *Cognitive models of memory*. Cambridge, MA: MIT Press.

Davis, E. 1998. The naïve physics perplex. *AI Magazine*, Winter, 51-79.

Falkenhainer, B., Forbus, K., & Gentner, D. 1989. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 41, 1-63.

Fobus, K. 2001. Exploring analogy in the large. In Gentner, D., Holyoak, K., & Kokinov, B. (Eds.) *The Analogical Mind*. Cambridge, MA: MIT Press.

Forbus, K., Gentner, D., & Law, K. 1994. MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19(2), 141-205.

Forbus, K., Mostek, T. and Ferguson, R. 2002. An analogy ontology for integrating analogical processing and first-principles reasoning. *Proceedings of 2002 Innovative Applications of Artificial Intelligence*. Menlo Park, CA: AAAI Press.

Gentner, D. & Markman, A. 1997. Structure mapping in analogy and similarity. *American Psychologist* 52, 45-56.

Goldman, A. 2000. Folk Psychology and Mental Concepts. *Protosociology* 14, 4-25.

Gopnik, A. & Meltzoff, A. 1997. *Words, thoughts, and theories*. Cambridge, Mass.: Bradford, MIT Press.

Gordon, A. 2004. Strategy representation: An analysis of planning knowledge. Mahwah, NJ: Lawrence Erlbaum Associates.

Gordon, A., Kazemzadeh, A., Nair, A., and Petrova, M. 2003. Recognizing Expressions of Commonsense Psychology in English Text. *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*.

Gordon, A. 2002. The Theory of Mind in Strategy Representations. *Proceedings of the Twenty-fourth Annual Meeting of the Cognitive Science Society (CogSci-2002)*, George Mason University, Aug 7-10. Mahwah, NJ: Lawrence Erlbaum Associates.

Gordon, A. and Hobbs, J. 2003. Coverage and Competency in Formal Theories: A Commonsense Theory of Memory. *Proceedings of the 2003 AAAI Spring Symposium on Logical Formalizations of Commonsense Reasoning*, Stanford University, March 24-26, 2003.

Gratch, J. & Marsella, S. 2004. A Domain-independent Framework for Modeling Emotion, *Journal of Cognitive Systems Research* 5(4), 269-306.

Hobbs, J. & Gordon, A. 2005. Toward a Large-Scale Formal Theory of Commonsense Psychology for Metacognition. *Proceedings of the 2005 AAAI Spring Symposium on Metacognitive Computing*. March 21-23, Stanford, CA.

Leake, D. 1995. Abduction, Experience, and Goals: A Model of Everyday Abductive Explanation. *The Journal of Experimental and Theoretical Artificial Intelligence* 7, 407-428.

Nichols, S. and Stich, S. 2002. *How to Read Your Own Mind: A Cognitive Theory of Self-Consciousness*. In Q. Smith and A. Jokic. (eds.) *Consciousness: New Philosophical Essays*. Oxford University Press.

Pryor, L. & Colllins, G. 1992. Reference Features as guides to reasoning about opportunities. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, IN. Hillsdale, NJ: Lawrence Erlbaum Associates.

Rattermann, M., Spector, L., & Grafman, J. 2001. Total and partial-order planning: Evidence from normal and prefrontally damaged populations. *Cognitive Science*, 25, 941-975.

Sanders, K. E. 1989. A logic for emotions: A basis for reasoning about commonsense psychological knowledge. In G. Olson & E. Smith (Eds.) *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, Ann

Arbor, Michigan. Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, R. 1982. Dynamic Memory: A theory of reminding and learning in computers and people. Cambridge: Cambridge University Press.

Schut, M., Wooldridge M. & Parsons S. 2004. The Theory and Practice of Intention Reconsideration. *Journal of Experimental and Theoretical Artificial Intelligence*, vol 16(4), 2004, pp. 261-293.

Sloman, S. & Lagnado, D. 2005. Do We “do”? *Cognitive Science* 29, 5-39.

Smedslund, J. 1997. The Structure of Psychological Common Sense. Mahwah, NJ: Lawrence Erlbaum.

Stein, L. 1997. Postmodular systems: Architectural principles for cognitive robotics. *Cybernetics and Systems* 28(6), 471-487.

Swanson, Reid & Gordon, Andrew 2005. Automated Commonsense Reasoning About Human Memory. 2005 AAAI Spring Symposium on Metacognitive Computing. March 21-23, Stanford, CA.

Wimmer, H. & Perner, J. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children’s understanding of deception. *Cognition* 13, 103-128.

Zachary, W., Ryder, J., & Hicinbothom, J. 1998. Cognitive task analysis and modeling of decision making in complex environments, In Cannon-Bowers, J. & Salas, E. (Eds) Making decisions under stress: Implications for individual and team training. (pp. 315-344). Washington, DC, US: American Psychological Association.